High End Computing FAQ

Q: Where can I find information regarding the latest changes on discover?

A: At " Latest Changes on Discover" and also in the following Q/As.

Q: Why do I receive the following error running my MPI job with Scali:

```
someuser@discover06:/home/someuser>
mpirun -np 16 ./a.out
Taking nodenames from
"/tmp/pbs_mpimach31996", number of nodes
specified by -np
/opt/scali/bin/mpimon -stdin all  ./a.out  --  borgg123 8 borgg124
8
--- mpimon --- Could not connect to mpid on borgg123, Connection
refused ---
```

A: The recent Discover SCU3 upgrade added nodes (also known as "Harpertown" nodes) that do not support Scali MPI. As such, you should modify your PBS request statement to add the "scali=true" option to guarantee your job can use the Scali MPI library. An example:

```
#PBS -l
select=4:ncpus=4,walltime=2:00:00
changes to:
#PBS -l select=4:ncpus=4:scali=true,walltime=2:00:00
```

The NCCS recommends that users migrate away from ScaliMPI toward IntelMPI or OpenMPI. When using IntelMPI or OpenMPI, there is no restriction on the nodes you use; you do not need to specify "scali=true" and your job may jump into the queue faster because there is no restriction on where it can run.

Q: When I log in, or submit a batch script, I see the error

```
ModuleCmd_Load.c(200):ERROR:105:
Unable to locate a modulefile

  for

'comp/intel-9.1.038'

ModuleCmd_Load.c(200):ERROR:105: Unable to locate a modulefile

  for

'mpi/scali-5.3'
```

A: During the latest upgrade, many older compilers were removed (including all Intel v8 compilers, and some Intel v9 and v10 compilers). For versioning reasons, the Scali MPI module was renamed from "mpi/scali-5.3" to "mpi/scali-5".

The module you use most commonly may no longer exist on the system. Don't worry though; in most cases, newer compiler versions are backward-compatible with user applications and your results should not vary. You can use the "module avail" command to see all the modules available on the upgraded system.

Q Why do I get

If there is a compelling reason, or bitwise-reproducability is not maintained with the newer compilers, please contact support@nccs.nasa.gov and we can discuss access to the compiler of your choice.

Q: Which Baselibs installation do I use?

A: If you used to use /usr/local/other/baselibs/v2_2rp2_213_9.1.042meta_new (a commonly used GEOS-compatible tag) you can now use:

```
/usr/local/other/baselibs/ESMF220rp2_NetCDF362b6_9.1.052
```

OR

```
/usr/local/other/baselibs/ESMF220rp2_NetCDF362b6_10.1.015
```

(Depending on the Intel compiler you have loaded). If you choose to use Intel 9.1.042, it is compatible with the Baselibs built with 9.1.052.

Q: Which MPI implementations are available on discover?

A: At present, two MPI implementations are available on discover. They are the Intel MPI and the Scali MPI. You can select either of these implementations by loading the appropriate module:

```
module load mpi/scali-5
```

OR

```
module load mpi/impi-3.1.038
```

NOTE: The Scali MPI works only on the current nodes of discover, that is, the dempsey and the woodcrest nodes. When the new harpertown nodes come on line, only the Intel MPI will work on those nodes as well as on the current nodes. In other words, only the Intel MPI will work across the entire system after the upcoming expansion of discover.

In the near future, other MPI implementations like OpenMPI and MPICH may be available and they will work across the entire system.

Q: Does one MPI implementation perform better than the other?

A: On discover, the Scali MPI and the Intel MPI have basically the same performance. On pleiades (the SGI system at NAS), some applications have shown that the Intel MPI and MVAPICH have similar performances but both are significantly faster than SGI's MPT.

Q: Which MPI implementations have C++ bindings?

A: Both the Intel MPI and SGI MPT have C++ bindings but the Scali MPI and MVAPICH do not. It is expected that OpenMPI and MPICH, when they get installed on discover, will have the bindings.

Q: Why do I get this error while using the Intel MPI on discover?

```
You can
't run mpdboot on ['borga128
'] version of python must be >= 2.4, >
current ['
']
```

A: The head node in your PBS allocation needs to ssh to each of the child nodes without a password to start up a mpd process (part of mpdboot) before the mpiexec is executed. Therefore you need to have passwordless SSH keys set up. Please check here for how to do this. Note that in this case both the local and the remote machines are one and the same, that is, discover.

Q: How do that read large data set using Intel compiler versions 10.1.011, 10.1.013, and 10.1.015?

The following code aborts with:

forrtl: severe (67): input statement requires too much data, unit 19, ...

for Intel compiler versions 10.1.011, 10.1.013, and 10.1.015. Works with **earlier** or **later** versions of the compiler.

program io_spr

parameter (ip1=655390720)

dimension x(ip1)

open (19, file='sfile.19', form='unformatted', status='unknown')

x = 0.0

write(19) x

rewind(19)

read(19) x

stop

end

A: This is the max record length feature. On the IBM it was a 2GB limit, On the Altix it was 4GB. In all cases you could write huge records past 4GB with no complaint, but you wrote garbage that couldn't be read, and you would get an error message like the one mentioned above.

The fix for the problem is in v10.1.018.